PREFERENCE-AWARE CONSTRAINED MULTI-OBJECTIVE BAYESIAN OPTIMIZATION

Alaleh Ahmadianshalchi, Syrine Belakaria, Janrdhan Rao Doppa

WASHINGTON STATE UNIVERSITY

Motivation

Using Bayesian Optimization (BO) to propose a solution for the expensive black-box multi-objective optimization problem with constraints when the domain practitioner has preferences over specific objectives.

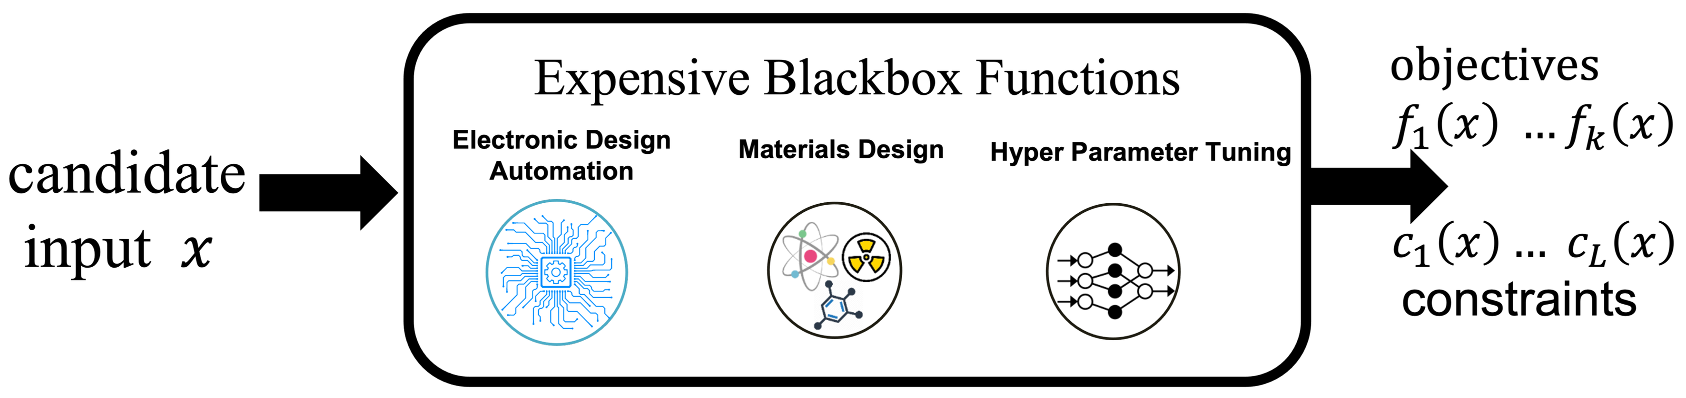

Drawbacks of Existing Methods

- Unable to handle preferences
- Unable to optimize expensive function in minimal iterations
- Unable to find feasible regions in the design space.

Preference-Aware Constrained Multi-Objective Bayesian Optimization (PAC-MOO)

- Selects the candidate input x for evaluation that maximizes the information gain about the optimal Pareto front $\boldsymbol{Y}^*$
- Equivalent to expected reduction in entropy over the Pareto front $\boldsymbol{Y}^*$
- Relies on a computationally cheap and low-dimensional $m.k \ll m.d$ distribution, where k is the number of objectives
- Includes a convex combination of the information gain about the output of all objectives and constraints.

BO for Multi-Objective Optimization

BO: A framework to **optimize expensive black-box functions** including three main elements:

- **Statistical models:** as a prior for the functions.
- **Acquisition function:** score the utility of evaluating input x
- **Optimization procedure:** select the best input for evaluation

Key Advantages of PAC-MOO

- Able to handle preferences over all black-box functions
- Able to find feasible regions in the input space
- Scalable for high-dimensions via output space entropy search
- Tight approximation with closed-form expression

PAC-MOO Acquisition Function

- The expensive black-box objective and constraints functions are defined as $F = \{f_1, \ldots, f_K\}$ and $C = \{c_1, \ldots, c_L\}$.
- S is the number of samples, $\gamma_s^i(x) = \frac{y_s^{i*} - \mu_i(x)}{\sigma_i(x)}$, where $i \in F \cup C$. Φ and ϕ are the p.d.f and c.d.f of a standard normal distribution, respectively.

$$AF(i,x) = \sum_{s=1}^{S} \frac{\gamma_s^i(x)\phi(\gamma_s^i(x))}{2\Phi(\gamma_s^i(x))} - ln\Phi(\gamma_s^i(x))$$

$$\alpha_{pref}(x) \approx \sum_{i \in F \cup C} p_i \times AF(i,x) \text{ s.t. } \sum_{i \in F \cup C} p_i = 1$$

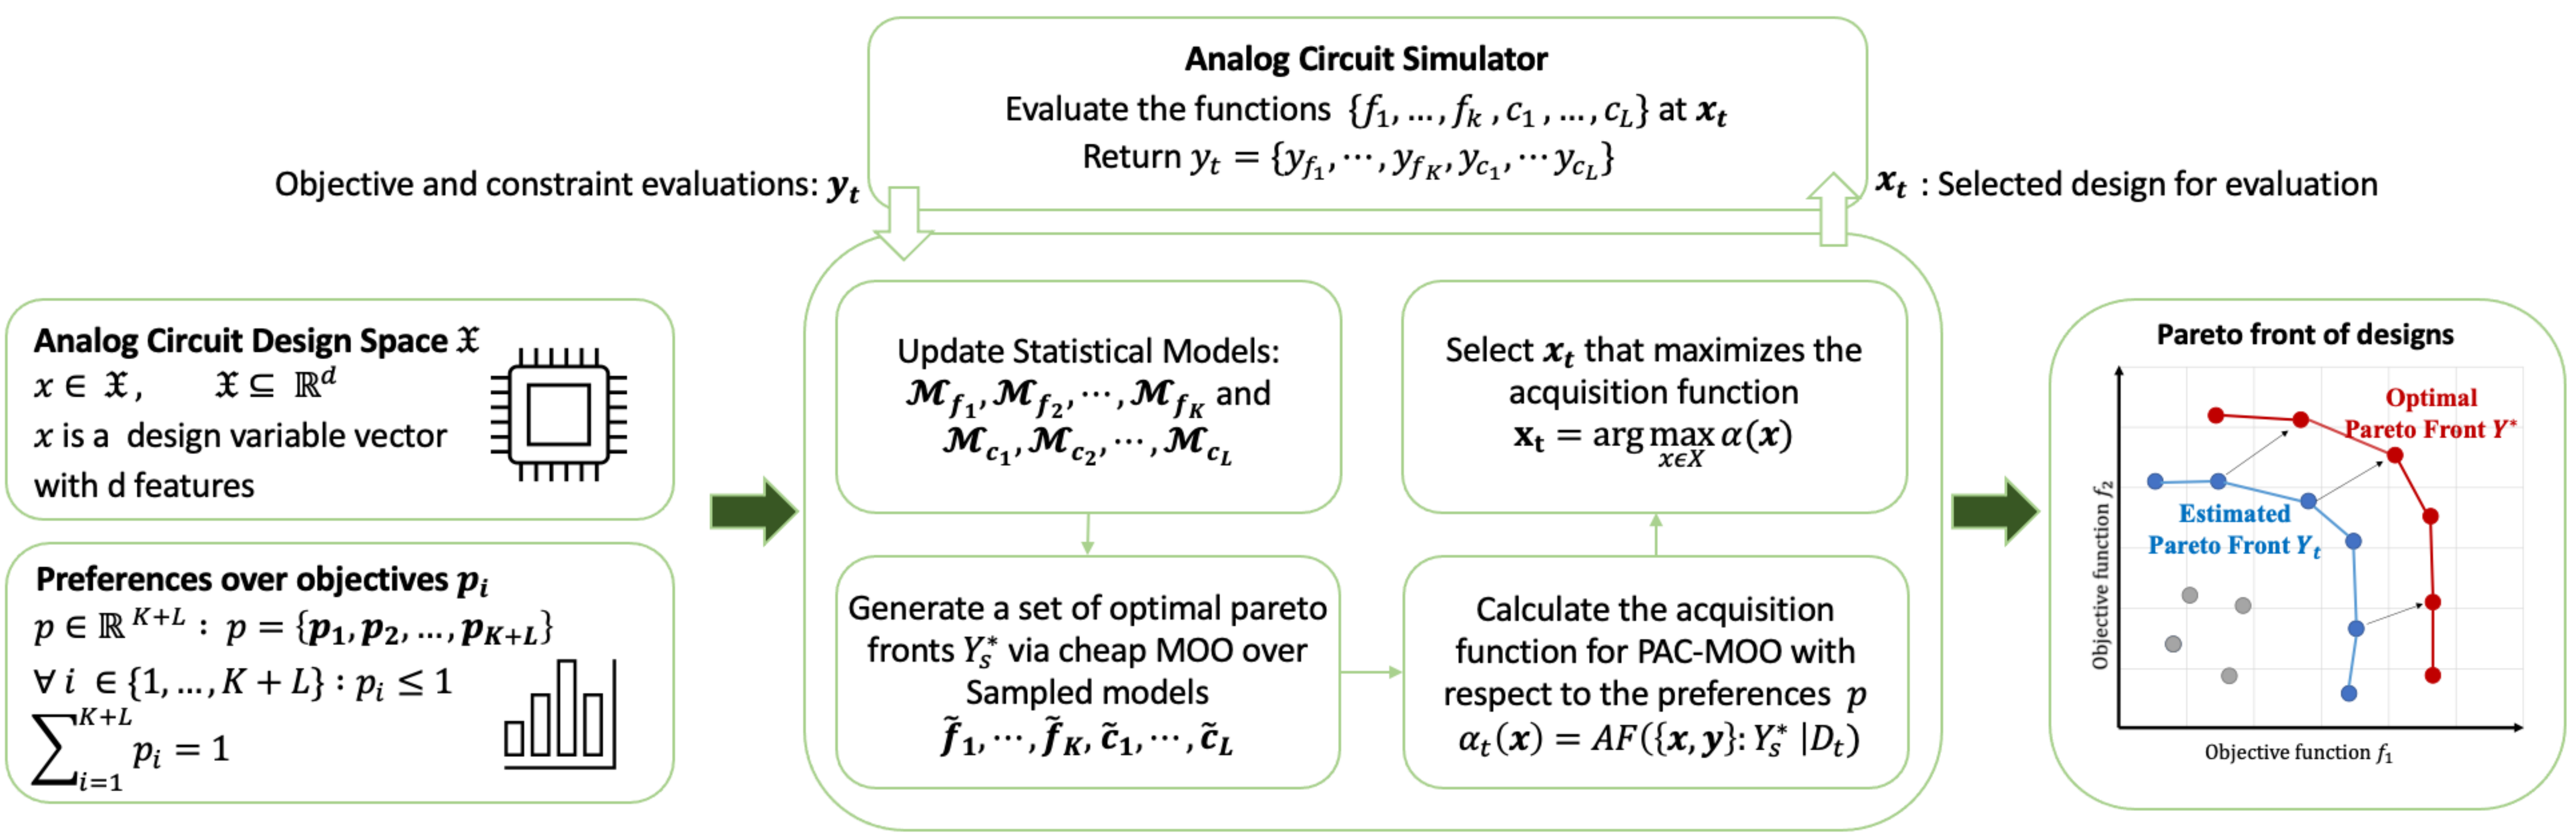

Results and Discussion

- We use two analog circuit design benchmarks named SCVR and HCR
- In PAC-MOO-0 the preferences are equal over all objectives and constraints. PAC-MOO-(1,2, and 3) assign increasingly high preferences to the efficiency objective.
- In all experiments, assigning higher preference to the efficiency objective, results in the algorithm finding feasible designs with considerably higher efficiencies.
- Assigning preferences does not significantly reduce the hypervolume performance of the multi-objective solver unless an unusually high preference is used